

# Virtual Interface Architecture

Bill Saphir

3/9/2000



## Commodity Supercomputing



- High volume microprocessors. Multiple vendors, leverages mass market for price
- High volume system boards. Multiple vendors, leverages mass market
- PCI peripherals
- Open source base. Leverages mass market, enables hardware portability
- Myrinet (Quadrics/T-Net?)
- Standards: C, Fortran, MPI.

## Commodity networks




- Fast Ethernet and Gigabit Ethernet limit scalability, applications.
  - Very fast networking requires hardware support these don't have.
  - Scalable topologies are necessary.
- PCI is a bottleneck (esp latency)
- HPC networks have been and still are special-purpose. SP switch; T3E; Myrinet; Quadrics.

## Infiniband



- Industry standard now targeted at 2001 (?)
- Solves PCI and fast networking problems (?)
- High volume, multiple vendors



## VIA as a software interface



- Standardized API
  - Corner cases defined
  - Well-defined error behavior
  - Robust protection mechanism
  - More applications than MPI
- VI Developer Forum working on VIPL 1.1, VIPL 2.0.
- 50,000+ lines of conformance test code

## VIA as a hardware interface



- Hardware understands VIA descriptors.
- Virtual memory-based protection. (NIC knows how to do address translations, etc.)
- Bypassing the operating system does not require fancy coprocessing on the NIC. (Application bypass is a different story).

## Scalability/Performance



- There are some theoretical arguments that VIA is not the right interface for HPC.
  - Some of these may prove to be correct
  - Some of these are NIH
- Open source and open metrics can resolve this

## Do we need more than MPI?



Scalable clusters require more than MPI

- Fast low-overhead monitoring
- Fast general purpose I/O
- Fast process startup
- Network-attached storage
- Commercial applications?

Some folks don't need/want MPI